

IN THE CLAIMS

This Listing of Claims will replace all prior versions, and listings, of claims in the application.

1-26. (Cancelled)

27. (Currently Amended) A method of managing a cluster of networked resources and resource groups using rule-based constraints in a scalable clustering environment, the method comprising the steps of:

building a globally optimal cluster configuration of said networked resources in accordance with said rule-based constraints and a current state of said resources, including identifying for each of the resources and resource groups an availability and quality of service, which are determined by dependencies among the resources and resource groups, resource equivalency, user preferences, constraints on the resources and network policies;

bringing said cluster of networked resources on-line in a systematic manner, given the current state of each of the resources and resource groups, and their dependencies, user preferences, constraints on the resources, and network policies,

with said cluster of networked resources on-line, determining dynamic dependencies of and configuration information about said cluster of networked resources (i) statically at said step of building and said step of bringing said cluster of networked services online and (ii) dynamically during cluster operation in accordance with said rule-based constraints,

supporting startup, operation and shutdown of said cluster of networked resources according to current policies, and system events, and said rule-based constraints;

separating said dependencies among resources and resource groups, user preferences, constraints among the resources, system events, and current policies into (i) a first, static rules based group and (ii) a second, dynamically changing events based group, wherein said first group captures the static resources, including, for each resource, a type and quality of the supporting resources needed to enable said each resource, wherein said step of separating is implemented according to said rule-based constraints; and

combining said first and second groups in a systematic manner only when needed to build the said globally optimal cluster configuration, and only when needed during operation to modify and realign the current state of said cluster to said globally optimal cluster configuration, or an alternative globally optimal in view of said current policies, said system events and said rule-based constraints; wherein:

the providing step includes providing a preprocessor module and a postprocessor module;

wherein the preprocessor module includes a preprocessor entry queue, the optimizer module includes an optimizer input queue, and the postprocessor module includes a postprocessor input queue;

creating a preprocessor task to implement a decision to reallocate a resource or a resource group, wherein said preprocessor task is provided to the entry queue of the preprocessor module; and

wherein said preprocessor task is an object having an entry method that, when invoked, results in the preprocessor task being executed, and execution of the preprocessor task results in either a postprocessor task being provided in the postprocessor input queue, an optimizer task being provided in the optimizer input queue, or both;

scheduling the postprocessor task by an invocation of the entry method associated with the postprocessor task;

scheduling the optimizer task by an invocation of the entry method associated with the optimizer task; and

executing the optimizer task results in a postprocessor task provided in the postprocessor input queue.

28-31. (Cancelled)